

Atty Docket No. 50253-217; P3709

APPLICATION FOR UNITED STATES LETTERS PATENT FOR

TECHNIQUES FOR PERMITTING ACCESS ACROSS A CONTEXT BARRIER ON A SMALL FOOTPRINT DEVICE USING AN ENTRY POINT OBJECT

INVENTORS:

Joshua Susser
216 Dorland Street
San Francisco, CA 94114

Citizenship: USA

Mitchel B. Butler
522 N. Cascade Terrace
Sunnyvale, CA 94087

Citizenship: USA

Andy Streich
693 Beach Park Boulevard
Foster City, CA 94404

Citizenship: USA

PREPARED BY:

David L. Stewart
McDERMOTT, WILL & EMERY
99 Canal Center Plaza, Suite 300
Alexandria, VA 22314
Telephone: (202) 756-8601
Facsimile: (202) 756-8699

50253-217P3709

TECHNIQUES FOR PERMITTING ACCESS ACROSS A
CONTEXT BARRIER ON A SMALL FOOTPRINT DEVICE
USING AN ENTRY POINT OBJECT

CROSS-REFERENCES TO RELATED APPLICATIONS

This application is related to U.S. Patent Application Serial Number 08/839,621 filed April 15, 1997, entitled "VIRTUAL MACHINE WITH SECURELY DISTRIBUTED 5 BYTE CODE VERIFICATION", in the name of inventors Moshe Levy and Judy Schwabe (Docket No. 50253-221/P3263), which application is incorporated herein by reference in its entirety.

This application is related to U.S. Patent 10 Application Serial Number _____ filed January 22, 1999, entitled "TECHNIQUES FOR IMPLEMENTING SECURITY ON A SMALL FOOTPRINT DEVICE USING A CONTEXT BARRIER", in the name of inventors Joshua Susser, Mitchel B. Butler, and Andy Streich, (Docket No. 50253-216/P3708), which 15 application is incorporated herein by reference in its entirety.

This application is related to U.S. Patent Application Serial Number _____ filed January 22, 1999, entitled "TECHNIQUES FOR PERMITTING ACCESS ACROSS A CONTEXT BARRIER ON A SMALL FOOTPRINT DEVICE USING RUN TIME ENVIRONMENT PRIVILEGES", in the name of inventors Joshua Susser, Mitchel B. Butler, and Andy Streich, (Docket No. 50253-218/P3710), which application is incorporated herein by reference in its entirety.

5 This application is related to U.S. Patent Application Serial Number _____ filed January 22, 10 1999, entitled "TECHNIQUES FOR PERMITTING ACCESS ACROSS A CONTEXT BARRIER IN A SMALL FOOTPRINT DEVICE USING GLOBAL DATA STRUCTURES", in the name of inventors Joshua Susser, Mitchel B. Butler, and Andy Streich, (Docket No. 15 50253-219/P3711), which application is incorporated herein by reference in its entirety.

10 This application is related to U.S. Patent Application Serial Number _____ filed January 22, 1999, entitled "TECHNIQUES FOR PERMITTING ACCESS ACROSS A CONTEXT BARRIER IN A SMALL FOOTPRINT USING SHARED 20 OBJECT INTERFACES", in the name of inventors Joshua Susser, Mitchel B. Butler, and Andy Streich, (Docket No. 50253-220/P3712), which application is incorporated herein by reference in its entirety.

BACKGROUND OF THE INVENTIONField of the Invention

The invention relates to computer security and more particularly to techniques for implementing a security on small footprint devices, such as smart cards.

Description of Related Art

A number of object oriented programming languages are well known in the art. Examples of these include the C++ language and the Smalltalk language.

Another such object oriented language is the JAVA™ language. This language is described in the book Java™ Language Specification, by James Gosling et al. and published by Addison-Wesley. This work is incorporated herein by reference in its entirety. The JAVA™ language is particularly well suited to run on a Java™ Virtual Machine. Such a machine is described in the book Java™ Virtual Machine Specification, by Tim Lindholm and Frank Yellin which is also published by Addison-Wesley and which is also incorporated herein by reference in its entirety.

A number of small footprint devices are also well known in the art. These include smart cards, cellular telephones, and various other small or miniature devices.

Smart cards are similar in size and shape to a credit card but contain, typically, data processing

capabilities within the card (e.g. a processor or logic performing processing functions) and a set of contacts through which programs, data and other communications with the smart card may be achieved. Typically, the set 5 of contacts includes a power source connection and a return as well as a clock input, a reset input and a data port through which data communications can be achieved.

Information can be written to a smart card and retrieved from a smart card using a card acceptance 10 device. A card acceptance device is typically a peripheral attached to a host computer and contains a card port, such as a slot, in to which a smart card can be inserted. Once inserted, contacts or brushes from a connector press against the surface connection area on 15 the smart card to provide power and to permit communications with the processor and memory typically found on a smart card.

Smart cards and card acceptance devices (CADs) are the subject of extensive standardization efforts, e.g. 20 ISO 7816.

The use of firewalls to separate authorized from unauthorized users is well known in the network environment. For example, such a firewall is disclosed in U.S. Patent Application Serial No. 09/203,719, filed 25 December 1, 1998 and entitled "AUTHENTICATED FIREWALL TUNNELLING FRAMEWORK" in the name of inventor David

Brownell (Docket No. 50435-023/P2789/TJC), which application is incorporated herein by reference in its entirety.

5 A subset of the full Java™ platform capabilities has been defined for small footprint devices, such as smart cards. This subset is called the Java Card™ platform. The uses of the Java Card™ platform are described in the following publications.

10 JAVA CARD™ 2.0 -- LANGUAGE SUBSET AND VIRTUAL MACHINE SPECIFICATION;

JAVA CARD™ 2.1 -- APPLICATION PROGRAMMING INTERFACES;

JAVA CARD™ 2.0 -- PROGRAMMING CONCEPTS;

JAVA CARD™ APPLET DEVELOPER'S GUIDE.

15 These publications are incorporated herein by reference in their entirety.

20 A working draft of ISO 7816 -- Part 11 has been circulated for comment. That draft specifies standards for permitting separate execution contexts to operate on a smart card. A copy of that working draft is hereby incorporated by reference in its entirety.

25 The notion of an execution context is well known in computer science. Generally speaking, the use of multiple execution contexts in a computing environment provides a way to separate or isolate different program modules or processes from one another, so that each can

operate without undue interference from the others. Interactions --if any-- between different contexts are deliberate rather than accidental, and are carefully controlled so as to preserve the integrity of each 5 context. An example of multiple contexts is seen in larger hardware devices, such as mainframes, where a plurality of virtual machines may be defined, each such virtual machine having its own execution context. Another example is seen in U.S. Patent No. 5,802,519 in 10 the name of inventor De Jong, which describes the use of multiple execution contexts on a smart card. It will be appreciated by those of skill in the art that a computing environment which provides multiple execution contexts also needs to provide a mechanism for associating any 15 given executing code with its corresponding context.

Also well known is the notion of a current context. Certain computing environments that support multiple contexts will, at any given time, treat one context in particular as an active focus of computation. The 20 context can be referred to as the "current context." When the current context changes, so that some other context becomes the current context, a "context switch" is said to occur. As will be appreciated by those of skill in the art, these computing environments provide 25 mechanisms for keeping track of which context is the current one and for facilitating context switching.

In the prior art, in the world of small footprint devices, and particularly in the world of smart cards, there was no inter-operation between contexts operating on the small footprint devices. Each context operated 5 totally separately and could operate or malfunction within its context space without affecting other applications or processes in a different context.

One layer of security protection utilized by the Java™ platform is commonly referred to as a sandbox 10 model. Untrusted code is placed into a "sandbox" where it can "play" safely without doing any damage to the "real world" or full Java™ environment. In such an environment, Java™ applets don't communicate, but each has its own name space.

15 Some smart card operating systems don't permit execution contexts to communicate directly, but do permit communications through an operating system, or through a server.

The Problems

20 A number of problems exist when trying to place computer programs and other information on a small footprint device. One of the compelling problems is the existence of very limited memory space. This requires often extraordinary efforts to provide needed 25 functionality within the memory space.

A second problem associated with small footprint devices is the fact that different small footprint device manufacturers can utilize different operating systems. As a result, applications developed for one operating system are not necessarily portable to small footprint devices manufactured by a different manufacturer.

If programs from more than one source of programs (manufacturer or vendor) are to be applied to a single small footprint device, security becomes a factor as one attempts to avoid corruption of existing programs and data when a new program is loaded on to the small footprint device. The same concern exists when one wishes to prevent a hacker or a malicious person from accessing programs and data.

It is clear that small footprint devices such as smart cards don't have the resources necessary to implement separate virtual machines. Nevertheless, it is desirable to maintain strict security between separate execution contexts.

In the past, security was provided by loading only applications from the same source or from a known trusted source onto a smart card or other small footprint device.

Accordingly, it would be desirable to allow object-oriented interaction between selected execution contexts only in safe ways via fast efficient peer to peer communications which do not impose undue burdens on the

programmer but facilitate dynamic loading of applets written at different times by untrusted sources.

SUMMARY OF THE INVENTION

5 The invention is directed to providing a context barrier (sometimes referred to as a firewall) for providing separation and isolation of one context from another and to provide controlled access across the barrier when that is needed.

10 In accordance with the invention, two execution contexts, e.g. each containing one or more applets, running in the same logical (i.e., virtual or real) machine, protected from each other, can share information in a controlled, secure way, using language mechanisms, such as object-oriented language mechanisms. Security 15 can be, for example, object by object. Thus, a method in a first execution context can access a first object A in a second execution context, but not a second object B in the second execution context on a selective basis.

20 In accordance with one exemplary embodiment, an enhanced Java™ Virtual Machine (VM) provides certain run-time checks of attempted access across execution contexts in the VM. Checks can be automatic by the VM or coded by the programmer with support from the VM. This can be done using language-level communication 25 mechanisms. In this way, one can express object access

5 across execution contexts in the same way as other object accesses using the language are made. These run-time checks provide a second dimension of defense/security beyond that which the Java™ language and platform already provide.

10 These mechanisms provide protection against, e.g., security holes due to programming bugs (such as declaring a datum "public" (global) when it shouldn't be accessible to all contexts). They also allow fine-grain control of sharing (such as selection of objects to share and applets to share to).

The invention is also directed to computer program products and carrier waves related to the other aspects of the invention.

15 The foregoing and other features, aspects and advantages of the present invention will become more apparent from the following detailed description of the present invention when taken in conjunction with the accompanying drawings.

20

BRIEF DESCRIPTION OF THE DRAWINGS

The features and advantages of the present invention will be apparent from the following description in which:

25 **Figure 1** is an illustration of a computer equipped with a card acceptance device and of a smart card for use with the card acceptance device.

Figure 2 is an illustration of a computer equipped with a card acceptance device connected to a network.

Figure 3 is an exemplary hardware architecture of a small footprint device, such as a smart card, of the prior art.

Figure 4 illustrates objects being accessed by principals as done in the prior art.

Figure 5 is an exemplary security model which can be used in explaining the various embodiments of the invention.

Figure 6 is a block diagram showing separation of execution contexts by a firewall or context barrier in accordance with one aspect of the invention.

Figure 7 is a representation of a software architecture useful in carrying out the invention.

Figure 8 is a flow chart of a security enforcement process implementing a firewall in accordance with one aspect of the invention.

Figure 9 is a block diagram showing object access across a firewall in accordance with one aspect of the invention.

Figure 10 is a block diagram showing cascaded object access across a firewall.

Figure 11 is a flow chart of a process for permitting access by a principal in one context across a firewall into another context.

Figure 12 is a block diagram illustrating the use of an entry point object to permit access across a firewall.

5 Figure 13 is a block diagram illustrating the use of a global data structure such as an array for access across a firewall.

Figure 14 is a block diagram illustrating the use of a supercontext to permit access across a firewall.

10 Figure 15 is a block diagram illustrating the use of shareable interface objects to permit access across a firewall.

Figure 16 is a flow chart of a security enforcement process permitting access across a firewall.

Figure 17 is the flow chart of Figure 16 showing details of block 1620.

15 Figure 18 is a flow chart showing an exemplary implementation of block 1629 of Figure 17.

NOTATIONS AND NOMENCLATURE

20 The detailed descriptions which follow may be presented in terms of program procedures executed on a computer or network of computers. These procedural descriptions and representations are the means used by those skilled in the art to most effectively convey the substance of their work to others skilled in the art.

25 A procedure is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired

result. These steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, 5 transferred, combined, compared, and otherwise manipulated. It proves convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like. It should be noted, however, that 10 all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities.

Further, the manipulations performed are often referred to in terms, such as adding or comparing, which 15 are commonly associated with mental operations performed by a human operator. No such capability of a human operator is necessary, or desirable in most cases, in any of the operations described herein which form part of the present invention; the operations are machine operations. 20 Useful machines for performing the operation of the present invention include general purpose digital computers or other computational devices.

The present invention also relates to apparatus for 25 performing these operations. This apparatus may be specially constructed for the required purpose or it may comprise a general purpose computer as selectively

activated or reconfigured by a computer program stored in the computer. The procedures presented herein are not inherently related to a particular computer or other apparatus. Various general purpose machines may be used 5 with programs written in accordance with the teachings herein, or it may prove more convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these machines will appear from the description given.

10

DETAILED DESCRIPTION

Attached as an Appendix to this specification is an unpublished draft of a document entitled JAVA CARD RUNTIME ENVIRONMENT 2.1 SPECIFICATION. This draft document, which provides further detailed description of 15 specific embodiments of the invention, is incorporated in its entirety as an integral part of the present specification.

Although the inventive techniques are described hereinafter in the context of a smart card example, the 20 example is merely illustrative and shouldn't limit the scope of the invention.

Figure 1 is an illustration of a computer 120 equipped with a card acceptance device 110 and a smart card 100 for use with the card acceptance device 110. In 25 operation, the smart card 100 is inserted into card

acceptance device 110 and power and data connections applied through a set of contacts 105 accessible at the surface of the smart card 100. When the card is inserted, mating contacts from the card acceptance device 110 interconnect with the surface contacts 105 to power-up the card and permit communications with the onboard processor and memory storage.

Figure 2 is an illustration of a computer equipped with a card acceptance device, such as 120 in Figure 1, connected to a network 200. Also connected to a network are a plurality of other computing devices, such as server 210. It is possible to load data and software onto a smart card over the network 200 using card equipped device 120. Downloads of this nature can include applets or other programs to be loaded onto a smart card as well as digital cash and other information used in accordance with a variety of electronic commerce and other applications. The instructions and data used to control processing elements of the card acceptance device and of the smart card may be stored in volatile or non-volatile memory or may be received directly over a communications link, e.g., as a carrier wave containing the instructions and/or data. Further, for example, the network can be a LAN or a WAN such as the Internet or other network.

Figure 3 is an exemplary hardware architecture of a small footprint device, such as a smart card, of the prior art. As shown in Figure 3, a processor 300 interconnects with primary storage 310 which may include 5 read only memory 315 and/or random access memory 316. The processor also connects with a secondary storage 320 such as EEPROM and with an input/output 330, such as a serial port. One can see the small footprint devices of this nature can be very simple.

Figure 4 illustrates objects being accessed by principals as done in the prior art. As shown in Figure 4, physical device 400, such as the small footprint device may have contained within it one or more processing machines (virtual or physical) which are 10 running an execution context 420. The execution context may be, for example, a context associated with a particular applet. One or more principals 430 (e.g., applets or applications) in the execution context may 15 seek to access other objects within the execution context. As long as the access occurs within the execution context, the accesses will be permitted and everything will function normally.

Figure 5 is an exemplary security model which can be 20 used in explaining the various embodiments of the invention. It is just one of many models which might be utilized but is a convenient model for this purpose. In 25

this model, a principal (sometimes called entity) 500 proposes to take an action 510 on an object, such as object 520. Security checks may be imposed on the principal, on the object, and/or on the action proposed 5 to be taken.

In Figure 5, two types of objects are shown on which action may be taken by a principal. These include data objects, (e.g. data1 and data2 (520, 520')) and entity 530. A principal may operate or attempt to operate on 10 any of these objects.

While data is passive, an entity 530 is active. The diagram line from Principal to an active entity is also labeled "action," but this could be a more sophisticated and arbitrarily complex action, such as making a function 15 or method call or sending a message as compared with action on a data object. As with data, a security check enforced by the operating system may use the identity of the principal, the identity of the entity, and/or the type of action. Furthermore, the entity, being active, 20 can perform its own additional security checks. These can be as arbitrarily complex as one desires, and can make use of the identity of the Principal, the identity of the entity itself, the action, and/or any other information that is available.

25 In an object-oriented system (such as the Java Card™ platform) "objects" are typically a combination of

data and entity. When a Principal tries to access a field of an object, this is a data access--a fairly simple action protected by a fairly simple security check. When a Principal tries to access a method of an 5 object, this is an entity access, which can be arbitrarily complex both in action and in security check.

Figure 6 is a block diagram showing separation of execution contexts by a firewall or context barrier in accordance with one aspect of the invention. The 10 physical device 400 and the machine 410 correspond to the same items shown in Figure 4. An execution context 420 shows one principal 430 attempting to access object 440 within the context. This access would normally succeed. However, execution context 420 also shows a principal 630 15 attempting to access object 640 of execution context 620, across a context barrier 600. Normally, this access would be prohibited as indicated by the ~~X~~ 636 where the action 635 crosses the context barrier 600.

Figure 7 is a representation of a software 20 architecture useful in carrying out the invention. This software architecture is shown as a run time environment 700. An operating system 710 for the small footprint device is commonly used. A virtual machine 720, in an exemplary embodiment of the invention, is implemented 25 over the operating system. The virtual machine could be a Java Card™ virtual machine or other virtual machine.

750 can associate objects with that context by recording the context's name in the object's header. Information in the object's header cannot be accessed by programs written in the object-oriented language, but is only 5 available to the virtual machine 720 itself. Alternately, the runtime system 740 can identify contexts by dividing the memory space into separate regions, each for a particular context, and correspondingly the object system 750 can associate objects with that context by 10 allocating the object's storage in that context's memory space.

Figure 8 is a flow chart of a security enforcement process implementing a context barrier in accordance with one aspect of the invention. When a principal invokes an 15 action on an object (800) a check is made to determine whether the object is within the context of the principal (810). If it is not, the action is disallowed (840). Otherwise, the action is permitted (830). This is the simplest form of context barrier or firewall. In one 20 specific embodiment the action is disallowed (840) by throwing a security exception if the object is outside of the namespace or the memory space of the context requesting access.

Figure 9 is a block diagram showing object access 25 across a firewall in accordance with one aspect of the invention. Figure 9 is substantially similar to Figure

6. However, **Figure 9** also shows principal 900 seeking to access object 910 in order to perform action 905 on the object 910. According to the invention, rather than having the access blocked by the firewall 600, in the way that action 635 is blocked, action 905 is permitted to occur across the firewall through access point 920 so that principal 900 can perform action 905 on object 910 notwithstanding the fact that the principal and the object are in different execution contexts. The 5 mechanisms behind access point 920 are described below with reference to **Figures 12-18**. Note that access point 920 can coexist with obstructed accesses such as **x 636**. Thus access point 920 provides fine-grain control of sharing (object by object security) across context 10 barrier 600.

15

When object access 900 is initiated, the current context setting is context 420. If the object 910 is a data object, the action 905 is a simple data access, and no code is executed in the second context 620. If the 20 object 910 is an entity object, and the action 905 results in that object's code being executed, that code is executed in the second context 620. To execute the code of object 910 in the correct context 620, the virtual machine 410 performs a context switch. The 25 context switch changes the current context setting to be context 620, and the previous value of the current

context setting is stored so that it can be restored later. From that point on code will execute in the new current context. When the action 905 completes, control is returned to the point following access 900. During 5 the return, the virtual machine 410 must restore the value of the current context setting to its previous value.

Figure 10 is a block diagram showing cascaded object accesses across a firewall. Figure 10 shows three 10 execution contexts, 1000, 1010 and 1020. Principal 1030 in execution context 1 seeks to invoke an action 1035 on object 1050 in execution context 2 and does so through access point 1070 in context barrier 600. Object 1050 in execution context 2 has an object access 1040 which seeks 15 to perform an action 1045 on the object 1060 in execution context 3. It achieves this by using access point 1080 in context barrier 600' separating execution contexts 2 and 3. Object 1050 in execution context 2 also has another object access 1090 which invokes an action 1095 20 on an object 1099 in the same execution context, that is, in execution context 2. Both actions 1035 and 1045 result in context switches as described in the explanation of Figure 9. But as action 1095 does not cross the context barrier, a context switch is not 25 required for its execution, and therefore does not occur.

Figure 11 is a flow chart of a process for permitting access by a principal in one context across a firewall into another context. There are essentially three steps to this process. In execution context 2, an object to be accessed is created and designated as shared (1100). In execution context 1, the principal obtains a reference to the object in execution context 2 (1110). The principal in execution context 1 then invokes an action upon the object designated as shared in context 2 (1120).

With respect to identifying or designating a created object as shareable as discussed in item 1100 of Figure 11, this can be done, in accordance with a specific embodiment of the invention, by including a shareable attribute in the header of an object's representation. Information in an object's header cannot be accessed by programs written in the object-oriented language, but is only available to the VM itself.

Obtaining a reference to an object in another context is a special case of accessing an object in another context. A mechanism that provides access to an object in another context can make other objects available also. For instance, invoking a method on an object in another context may return a reference to a second object in a different context. An additional mechanism is required to allow an initial reference to an

object in a different context to be obtained. In a specific embodiment, references to certain well-known entry point objects can be obtained using a public API. Once the initial reference to an object in a different 5 context is obtained, further references can be obtained from that object, and so on.

There are four general approaches to obtaining information across a context barrier in accordance with the invention. These approaches can be utilized 10 individually or in combination in order to access an object across a context barrier or to obtain a reference of an object to be accessed across a context barrier (1110). These approaches are described in **Figures 12-18**.

Figure 12 is a block diagram illustrating the use of 15 entry point objects to permit access across a context barrier. As shown in Figure 12, some object 1200 in context 770 (context 1) desires access to information in supercontext 760. In the specific embodiment, a supercontext 760 contains at least one entry point object 20 1210. The entry point object 1210 can be published as part of a public API, or can be made available indirectly through a published API (e.g., in accordance with the mechanisms described previously with reference to Figure 11), so that each context subordinate to the supercontext 25 may communicate with the entry point object of the supercontext. (It will be appreciated that in other

embodiments, entry point objects may be housed by a context other than the supercontext.)

Figure 13 is a block diagram illustrating the use of global data structures to permit access across a firewall. In this approach, supercontext 760 creates a global data structure such as a global array. In the specific embodiment supercontext 760 is the only context permitted to create such a global data structure. (It will be appreciated that in other embodiments, global data may be housed by a context other than the supercontext.) By virtue of its global status, each of the contexts 770 and 780 may read and write to the global data structure. Thus, information written into the global data structure by one context can be read by another context. For example, this mechanism can be used to pass binary data or references to objects between contexts.

Figure 14 is a block diagram illustrating the use of supercontext privileges to permit access across a context barrier. In Figure 14, an object in supercontext 760 seeks access to context 780 across the context barrier separating the two. Supercontext 760 can invoke any of the methods of context 780 and can access any of the data contained within context 780, by virtue of the privileges associated with the supercontext.

Figure 15 is a block diagram illustrating the use of shareable interface objects to permit access across a firewall. A shareable interface defines a set of shareable interface methods. A shareable interface object is an object that implements at least the set of methods defined in a shareable interface. In Figure 15, object 1210 in context 2 (780) is a shareable interface object. An object access 1200 in another context 770 can invoke any of the shareable interface methods on the object 1210 if the principal of the object access 1200 is authorized to do so by the object 1210 itself. This authorization is further discussed with reference to Figure 18 below.

It will be appreciated that a virtual machine consistent with the invention provides functionality beyond that of earlier virtual machines, such as the virtual machine described in the Java™ Virtual Machine Specification. In particular, consistently with the invention, the virtual machine provides functionality to implement or to facilitate a security enforcement process that permits access across a firewall. This process is described next with reference to Figures 16-18. Note that it is applicable to any approach for providing access across the firewall, including but not limited to the four approaches described with reference to Figures 12-15 above.

Figure 16 is a flow chart of a security enforcement process permitting access across a firewall. When a principal attempts to invoke action on an object 1600, a check is made to determine if the object is within the context of the principal (1610). If it is, (1610-Y), the action is permitted (1630). If it is not, (1610-N), a check is made to see if the action by the principal is permitted on the object (1620). If it is, (1620-Y), the action is permitted (1630). If it is not, (1620-N), the action is disallowed. In the specific embodiment a security exception is thrown (1640).

Figure 17 is the flow chart of Figure 16 showing further details of block 1620. If the object is not within the context of the principal (1610-N), a plurality of tests, 1621, 1622, 1623... 1629 are undertaken to see if the action by the principal is permitted on the object. These tests can be done by the virtual machine alone or by the virtual machine plus the object, in a virtual machine object oriented implementation. If any of the tests results in a pass, the action is permitted (1630). However, if all tests result in a negative determination (162X--No), the action will be disallowed. In a specific embodiment, a security exception will be thrown (1640). These tests relate to the permitted access discussed in conjunction with Figures 12-15.

Figure 18 is a flow chart showing an exemplary implementation of block 1629 of Figure 17 for use with access method described in Figure 15. In a test, such as 829 or 1629, a virtual machine checks if the object is a 5 shared object 1810. If it is not (1810-No), the test will fail. However, if it is (1810-Yes), the virtual machine will invoke the method A on object O (1820). If the method A on object O determines that the principal is authorized (1830), the test will be passed (1840) and 10 access permitted. Otherwise, the test will fail (1850). This allows the authorization text to be programmed into the code of the object itself.

Although the invention has been illustrated with respect to a smart card implementation, the invention 15 applies to other devices with a small footprint, not just to smart cards. Devices with a small footprint are generally considered to be those that are restricted or limited in memory or in computing power or speed. Such small footprint devices may include boundary scan 20 devices, field programmable devices, pagers and cellular phones among many others.

In general, small footprint devices are resource constrained computational devices and systems where 25 secure interoperation of execution contexts is a concern. Such small devices impose constraints on the implementation of security measures because of their

limited resources. Because of resource constraints, in a virtual machine implementation, a single virtual or physical machine must be used as opposed to multiple virtual machines.

5 The invention may also be applied to devices with larger footprints where the characteristics of the invention may prove beneficial. For example, the invention may prove advantageous when using servlets if there is object sharing between them. Even some desktop
10 systems may profitably utilize the techniques of the invention.

15 While the Java™ language and platform are suitable for the invention, any language or platform having certain characteristics would be well suited for implementing the invention. These characteristics include type safety, pointer safety, object-oriented, dynamically linked, and virtual-machine based. Not all of these characteristics need to be present in a particular implementation. In some embodiments,
20 languages or platforms lacking one or more of these characteristics may be utilized. A "virtual machine" could be implemented either in bits (virtual machine) or in silicon (real/physical machines).

25 Although the invention has been illustrated showing object by object security, other approaches, such as class by class security could be utilized.

Although the present invention has been described and illustrated in detail, it is clearly understood that the same is by way of illustration and example only and is not to be taken by way of limitation, the spirit and scope of the present invention being limited only by the terms of the appended claims and their equivalents.

EV 263 600 854 US

Docket No. SUN-P3709CNT

Copyright © 1994 Sun Microsystems, Inc.

901 San Antonio Road, Palo Alto, CA 94301 USA

All rights reserved. Copyright in this document is owned by Sun Microsystems, Inc.

Sun Microsystems, Inc. (SUN) hereby grants to you at no charge a non-exclusive, nontransferrable, worldwide, limited license (without due right to sublicense) under SUN's intellectual property rights that are essential to practice the Java™ Class 14 Java™ Environment Environment (JCER) 2.1 Specification ("Specification") to use the Specification for internal evaluation purposes only. Other than this limited license, you acquire a copyright, title, or interest in or to the Specification and you shall have no right to use the Specification for productive or commercial use.

RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure by the U.S. Government is subject to revision of FAR 12.227-14(a)(2)(v)(B) and FAR 32.227-19(d)(1), as DIFAR 231.227-2013 (b)(d)(3) and DIFAR 221.7203-1(a).

SUN MAKES NO REPRESENTATIONS OR WARRANTIES ABOUT THE SUITABILITY OF THE SOFTWARE, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. SUN SHALL NOT BE LIABLE FOR ANY DAMAGES SUFFERED BY LICENSEE AS A RESULT OF USING, MODIFYING OR DISTRIBUTING THIS SOFTWARE OR ITS DERIVATIVES.

TRADEMARKS

SUN, the SUN logo, SUN Microsystems, Java™, Java Card, JavaView, Visual Java, Solaris, NeXT, NeXT, Inc., NeXT, NFS, QNC, QNC, PC-NFS, OpenWindows, PC-NFS, Embedded Java, PC-NFS, Sun Workstation, The Network is The Computer, ToolTalk, Ultra, Ultracomputing, UltraServer, White Tie Network, In Group, Sun Workshop, KView, Java Workshop, Java, and Visual Java are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE FAKTIONALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THIS PRODUCT(S) AND/OR THIS PROGRAM(S) DISCRETELY IN THIS PUBLICATION AT ANY TIME.

APPENDIX

Java™ Card™ Runtime Environment (JCER)

2.1 Specification

Draft 2

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94301 USA
(650) 960-4100

Draft 2, December 14, 1998

Java™ Card™ Runtime Environment (ICRE) 2.1 Specification**Contents****Java™ Card™ Runtime Environment (ICRE) 2.1 Specification**

1. Introduction	1-1
2. Structure of the Java Card Virtual Machine	2-1
3. Java Card API and Behavior	3-1
3.1 The Adapted Java API	3-1
3.2 The Method Select	3-2
3.3 The Method process	3-2
3.4 The Method doSelect	3-3
3.5 Power Loss and Reset	3-3
4. Transient Objects	4-1
4.1 Objects That Aren't Transient Objects	4-2
5. Selection	5-1
5.1 The Default Apject	5-1
5.2 SELECT Command Processing	5-2
5.3 New SELECT Command Processing	5-3
6. Apject Selection and Object Sharing	6-1
6.1 Apject Flushing	6-1
6.1.1 Contexts and Context Switching	6-1
7. Transactions and Atomality	7-1
7.1 Atomality	7-1
7.2 Transaction	7-1
7.3 Transaction Duration	7-2
7.4 Nested Transactions	7-2
7.5 Tear or Reset Transaction Failure	7-2
7.6 Aborting a Transaction	7-3
7.7 Financial Objects	7-3
7.8 Commid Capacity	7-3
7.9 API Topics	7-4
7.10 The APDU Class	7-4
7.11 T=0 specific for outgoing data transfers	8-4
Copyright © December 14, 1998 Sun Microsystems, Inc.	ii
iv Copyright © December 14, 1998 Sun Microsystems, Inc.	iv

Java™ Card™ Runtime Environment (JCРЕ) 2.1 Specification

1. Java™ Card™ Runtime Environment (JCРЕ) 2.1 Specification	1
1.1.1 T=1 Specifier for Protecting Applets	8-1
1.1.2 The Security and Cryptography API	8-4
1.1.3 JCРЕ API Class	8-5
2. Virtual Machine Topics	9-1
2.1 Resource Failures	9-1
10. Applet Interface	10-1
10.1 The Hostlet	10-4
10.1.1 Hostlet Implementation	10-1
10.1.2 Hostlet API	10-2
10.1.3 Hostlet API URLs	10-2
10.1.4 Hostlet Behavior	10-2
10.1.5 Hostlet Privileges	10-3
10.2 The Newly Parallel API	10-3
10.2.1 Parallelization Parameters	10-3
11. API Conventions	1

Java Card™ technology combines a portion of the Java programming language with a runtime environment optimized for smart cards and other small-memory embedded devices. The goal of Java Card technology is to bring many of the benefits of Java software programming to the resource-constrained world of smart cards.

This document is a specification of the Java Card Runtime Environment (JCРЕ) 2.1. A vendor of Java Card-enabled device provides an implementation of the JCРЕ. A JCРЕ implementation within the context of this specification refers to a vendor's implementation of the Java Card Virtual Machine (VM), the Java Card Application Programming Interface (API), or other components, based on the Java Card technology specifications. A *Reference Implementation* is an implementation of the Java Card platform produced by Sun Microsystems, Inc. Applications for the Java Card platform are referred to as Java Card applets.

Who Should Use This Specification?

This specification is intended to enable JCРЕ implementers to create and implement, develop, test, and maintain the Java Card technology specification, or in creating an extension to the Java Card Runtime Environment (JCРЕ). This specification is also intended for Java Card applets developers to help them gain a greater understanding of the Java Card technology specification.

Before You Read This Specification

Before reading this guide, you should be familiar with the Java programming language, the Java Card technology specification, and smart card technology. A good resource for becoming familiar with Java Card technology and Java Card technology is the Sun Microsystems, Inc. website, located at: <http://java.sun.com>.

How This Specification Is Organized

Chapter 1, "The Scope and Responsibilities of the JCРЕ," gives an overview of the services required of JCРЕ implementations.

Chapter 2, "Lifetime of the Virtual Machine," defines the lifetime of the Virtual Machine.

Java™ Card System Environmental (JCRE) 2.1 Specification

- Chapter 3, "Applié Lifelines," defines the lifelines of an applicé.
- Chapter 4, "Virtual Objects," provides an overview of virtual objects.
- Chapter 5, "Safeties," describes how the JCRE handles applicé exceptions.
- Chapter 6, "Applié Isolation and Object Sharing," describes applicé isolation and object sharing.
- Chapter 7, "Transactions and Atomicity," provides an overview of atomicity during transactions.
- Chapter 8, "API Topics," describes API functionality required of a JCRE but not explicitly specified in the Java Card 2 API Specification.
- Chapter 9, "Virtual Machine Topics," describes virtual machine specifics.
- Chapter 10, "Applié Initiators," provides an overview of the Applié Initiator.
- Chapter 11, "API Constants," provides the numeric value of constants that are not specified in the Java Card API 2.1 Specification.

Glossary is a list of words and their definitions to assist you in using this book.

Related Documents and Publications

References to various documents or products are made in this manual. You should have the following documents available:

- *Java Card 2 API Draft 2 Specification*, Sun Microsystems, Inc.
- *Java Card 2.0 Language Specification and Virtual Machine Specification, Chapter 14, 1997, Revision 1.0 Final*, Sun Microsystems, Inc.
- *Java Card API Developer's Guide*, Sun Microsystems, Inc.
- *The Java Language Specification by James Gosling, Bill Joy, and Guy L. Steele, Addison-Wesley, 1996, ISBN 0-201-63351-1*.
- *The Java Virtual Machine Specification (Java 2 Edition)* by Tim Lindholm and Frank Yellin, Addison-Wesley, 1996, ISBN 0-201-63451-X.
- *The Java Class Library: An International Reference (Java 2 Edition)* by Patrik Chen and Rosanna Lee, Addison-Wesley, two volumes, ISBN-0-201-310022 and 0201310021.
- *ISO 7816 Specification Part 1-6*.
- *EMV™ Integrated Circuit Card Specification for Payment Systems*.

EV 263 600 854 US

Docket No.SUN-P3709CNT

2. Lifetime of the Java Card Virtual Machine

In a PC or workstation, the Java Virtual Machine runs as an operating system process. When the OS process is terminated, the Java applications and their objects are automatically destroyed.

In Java Card technology the execution lifetime of the Virtual Machine (VM) is the lifetime of the card. Most of the information stored on a card shall be present even when power is removed from the card. Persistent memory (such as EEPROM) enables a card to store information when power is removed. Since the VM and the objects created on the card are used to represent application information that is persistent, the Java Card VM appears to run forever. When power is removed, the VM stops only temporarily. When the card is next reset, the VM starts up again and resumes its persistent object heap from persistent storage.

Aside from its persistent nature, the Java Card Virtual Machine is just like the Java Virtual Machine.

The card initialization phase is the time after masking, and prior to the start of card personalization and issuance. At the time of card initialization, the JCРЕ is initialized. The framework objects created by the JCРЕ serve for the lifetime of the Virtual Machine. Because the execution lifetime of the Virtual Machine and the JCРЕ framework span CAD sessions of the card, the lifetime of objects created by objects will also span CAD sessions. (CAD means Card Acceptance Device, or card reader. Card sessions are those periods when the card is inserted in the CAD, powered up, and exchanging streams of APDU's with the CAD. The card session ends when the card is removed from the CAD.) Objects that have this property are called persistent objects.

The JCРЕ implementation shall make an object persistent when:

- The Applet.register method is called. The JCРЕ stores a reference to the instance of the applet object.
- The JCРЕ implementation shall ensure that instances of class applet are persistent.
- A reference to an object is stored in a field of any other persistent object or in a class's static field. This requirement stems from the need to preserve the integrity of the JCРЕ's internal data structures.

EV 263 600 854 US

Java Card™ Runtime Environment (JCRE) 2.1 Specification

Java in CCard Runtime Environment (JCRE) 2.1 Specification

3.2 The Method select

Appli's remain in a suspended state until they are explicitly selected. Selection occurs when the JCRC receives a SELECT APDU in which the source data matches the AID of the Appti. Selection causes an update to become the currently selected Appti.

Prior to calling SELECT, the JCRC shall deactivate the previously selected Appti. The JCRC indicates this to the Appti by invoking the Appti's deactivate method.

The JCRC deactivates the Appti or selection by invoking its select method.

The Appti may decline to be activated by returning false from the select method or by throwing an exception. If the Appti returns true, the actual SELECT APDU command is supplied to the Appti in the subsequent call to its process method, so that the Appti can examine the APDU's contents. The Appti can process the SELECT APDU command exactly like it processes any other APDU command. It can expand to process the SELECT APDU with data (see the process method for details), or it can flag errors by throwing an exception with the appropriate SIV (return status word). The SIV and update response data are returned to the CAd.

The Appti's deactivate method shall return true when called during the select method. The Appti's deactivate method will continue to return true during the deactivate method, which is called to process the SELECT APDU command.

If the Appti declines to be activated, the JCRC will return an APDU response status word of 100..00..APDU_E_APDU_CANT_BE_ACTIVATED to the CAd. Upon selection failure, the JCRC status is set to indicate that no Appti is selected.

After successful selection, all subsequent APDUs are delivered to the currently selected Appti via the process method.

Docket No. SUN-P3799CNT

3. Java Card Applet Lifetime

For the purposes of this specification, a Java Card Appti's lifetime begins at the point that it has been permanently loaded into card memory, like A, and otherwise prepared for execution. It is the responsibility of this specification, Appti's refers to an Appti written for the Java Card platform). Appti's establish a relationship with the Appti's, a Appti's include code for the lifetime of the card. The JCRC interacts with the Appti's public methods (methods, deactivate, deactivate, deactivate, and process). An Appti shall implement the static methods in Table 11 included. If the Appti's static methods are implemented, the Appti's cannot be activated or initialized. A JCRC implementation shall call an Appti's deactivate, deactivate, and process methods as described below.

When the Appti is initialized and the smart card, the Appti instantiates a method is called *init* by the JCRC for each Appti instance present. The JCRC shall not call this Appti's constructor directly.

3.1 The Method initAll

When initAll is called, one object of the Appti's *initAll* method within the Appti is to create an instance of the Appti's class, and to register the instance. All other objects that the Appti will need during its lifetime can be created as it needs. Any other preparations necessary for the Appti to be selected and accessed by a CAd also can be done as is feasible. The *initAll* method initializes all parameters from the constants of the following two array parameters.

Typically, an Appti creates various objects, initialized them with predefined values, sets some internal Appti variables, and calls the Appti's *register* method to specify the AID (Appti Identifier) as defined in ISO 18045-1 to be used to select it. This initialization is considered successful when the call to the Appti's *register* method completes without an exception. The initialization is deemed unsuccessful if the *initAll* method does not call the Appti's *register* method, or if an exception is thrown from within the *initAll* method prior to the Appti's *register* method or if the Appti's *register* method or any other method on the Appti's *initAll* method performs an cleanup when it regains control. That is, all persistent objects shall be returned to the Mac they had prior to calling this *initAll* method. If the *initAll* method is successful, the JCRC can mark the Appti as available for selection.

At any time during processing, the Appti may throw an exception with an update status word, in which the JCRC catches the exception and return the SIV to the CAd.

If any other exception is thrown during processing, the JCRC catches the exception and returns the status word 1507016. In addition to the CAd.

Java™ Card™ Runtime Environment (JCRE) 2.1 Specification**3.4 The Method deselect**

When the SCB receives a SELECT APDU command in which the name matches the AID of an applet, the JCRE calls the DESELECT method of the currently selected applet. This allows the applet to perform any cleaning operations that may be required in order to allow reuse when another applet is selected.

The `Applet.deSELECT()` method shall return `false` when called during the `deselect` method. Execution shall be caught by the `deselect` method as caught by the JCRE, but the applet is dissociated.

3.5 Power Loss and Reset

Power loss occurs when the card is withdrawn from the CAD or if there is some other mechanical or electrical failure. When power is supplied to the card and an `Card Reset` (warm or cold) the JCRE shall ensure that:

- Transient data is reset to the default value.
- The transaction in progress, if any, whose power was lost (or reset occurred) is aborted.
- The applet that was selected when power was lost (or reset occurred) becomes logically deselected. (In this case the deselect method is not called.)
- If the JCRE implements default applet selection (see paragraph 3.1), the default applet is selected as the currently selected applet and that the default applet's `onReset` method is called. Otherwise, the JCRE is in a state to handle the new applet is selected.

4.1 Events That Clear Transient Objects

- Applets, sourcecode, require objects that contain language (transient) data that need not be persistent across a CARD session. Java Card does not support the do yo beyond a reuse of one. (However, Java Card technology provides methods to create transient arrays with primitive components or references to objects.)
- The term "Transient object" is ambiguous. It can be interpreted to mean that the object itself is transient. However, only the contents of the fields of the object (except for the length field) have a transient nature. As with any other object in the Java programming language, transient objects within the Java Card platform exist as long as they are referenced from:

 - The stack
 - Local variables
 - A class static field
 - A field in another existing object
 - A transient object within the Java Card platform has the following required behavior:

- The fields of a transient object shall be consistent to the field's default value (zero, false, or null) at the occurrence of certain events (see below).
- For security reasons, the fields of a transient object shall never be stored in a "transient memory technology". Using current sunmi card technology as an example, the contents of transient objects can be stored in RAM, but never in EEPROM. The purpose of this requirement is to allow transient objects to be stored in static sections keys.
- Writing to the fields of a transient object shall not have a performance penalty (using current sunmi card technology). As an example, the transient transient objects can be stored in RAM, while the transient objects of non-transient objects can be stored in EEPROM. (Typically, RAM technology has a much faster write speed than EEPROM.)
- Writing to the fields of a transient object shall not be affected by "Transactions." That is, an above operation will never cause a field in a transient object to be restored to a previous value. This behavior makes transient objects ideal for small amounts of temporary update data that is frequently modified, but that need not be preserved across CAPI or select sessions.

Transient objects are used for establishing states that shall be preserved across card resets. When a transient object is selected, one or two cards are specified that cause its fields to be cleared. CLEAR_ON_RESET transient objects are used for establishing states that shall be preserved across object relocations, but not across card resets. CLEAR_ON_DESTROY transient objects are used for establishing states that shall be preserved while an object is selected, but not across object relocations or card resets.

Details of the two clear events are as follows:

- CLEAR_ON_RESET**—the object's fields are cleared when the card is reset. When a card is powered on, it also causes a card reset.
- Notes**—It is not necessary to clear the fields of transient objects before power is restored from a card. However, it is a good way to guarantee that the previous contents of said fields cannot be recovered once power is lost.
- CLEAR_ON_DESTROY**—the object's fields are cleared whenever any update is deselected. Because a card can rapidly deactivate the currently selected object, the fields of CLEAR_ON_DESTROY objects are also cleared by the main events specified for CLEAR_ON_RESET.
- The currently selected object is explicitly deselected (via `destroy`) only when a `SELECT` command is processed. The currently selected object is deselected and then the fields of all CLEAR_ON_DESTROY transient objects are cleared regardless of whether the SELECT command:

 - fails to select an object.
 - selects a different object.
 - receives the same object.

EV 263 600 854 US

Java™ Card™ Runtime Environment (JCRE) 2.1 Specification

Java™ Card™ Runtime Environment (JCRE) 2.1 Specification

5. Selection

5.1 The Default Applet

Card receive requests for service from the CAD in the form of APDUs. The SELECT APDU is used by the JCRC to designate a currently selected applet. Once selected, an applet receives all subsequent APDUs until the applet becomes deselected.

There is no currently selected applet when either of the following occurs:

- The card has reset and no applet has been pre-designated as the default applet.
- A SELECT command fails when attempting to select an applet.

5.2 SELECT Command Processing

The SELECT APDU command is used to select an applet. Its behavior is:

1. The SELECT APDU is always processed by the JCRC regardless of state. If any, apply it to active.
2. The JCRC sends the internal table for a matching AID. The JCRC shall support selecting an applet whose the full AID is present in the SEL_ECT command.

JCRE implementations are free to cause the JCRC to support other selection methods. An example of that is selecting via partial AID match as specified in ISO 7816-4. The specific requirements are as follows:

Note - An asterisk indicates binary bit numbering as in ISO 7816. Most significant bit = bit 0. Least significant bit = bit 1.

- a) Applet SEL_ECT command uses CLA=0x00, INS=0x04.
- b) Applet SEL_ECT command uses "3" selection by D4 name". Therefore, P1=0x04.
- c) Any other value of P1 implies that is not an applet selected. The AID is placed by the currently selected applet.
- d) JCRC shall support exact D4 name (AID) selection to 0x27-0x40000 to 0x0. (0x4000 are don't care.)
- e) All other partial D4 name SEL_ECT options (b7, b8*) are JCRC implementation dependent.
- f) All bits marked in formation option codes (b4, b5*) shall be supported by the JCRC and interpreted and processed by the applet.

Docket No. SUN-P3709GNT

3. If no AID match is found:
 - a. Match to an currently selected applet, the JCRC responds to the SEL_ECT command with status code 0x0999 (SV_APPLET_SEL_ECT_FAILED).
 - b. Otherwise, the SEL_ECT command is forwarded to the currently selected applet's process method.
4. If a matching AID is found, the JCRC prepares to select the new applet. If there is an currently selected applet, it is deselected via a call to its deactivate() method. A context switch into the deselected applet's context occurs at this point. The JCRC sends the SEL_ECT AID to the currently selected applet.
5. The JCRC sets the new currently selected applet. The new applet is selected via a call to its select() method, and a context switch into the new applet's context occurs.

- a. If the applet's select() method returns an exception or returns false, then JCRC state is set to 0x0999 (SV_APPLET_SEL_ECT_FAILED).
- b. The new currently selected applet's process method is then called with the SELECT_APDU as an input parameter. A context switch into the applet's context occurs.

Notes -

EV 263 600 854 US

Docket No.SUN-P3709CNT

Java™ Card™ Business Environment (JCRE) 2.1 Specification

If there is no matching AID, the SELECT command is forwarded to the currently selected applet (if any) for processing as a normal applied AIDN command.

If there is a matching AID and the SELECT command is, the JCRE always enters the state where no applet is selected.

If the matching AID is the same as the currently selected applet, the JCRE will skip through the process of deselecting the applets and then selecting in. Re-selection could fail, leaving the card in a state where no applet is selected.

5.3 Non-SELECT Command Processing

When a non-SELECT AIDN is received and there is no currently selected applet, the JCRE shall respond to the AIDN with status code 0x09 (SW_APP_ID_SW_ECT_FALIED).

When a non-SELECT AIDN is received and there is a currently selected applet, the JCRE shall respond to the process switch of the currently selected applet passing the AIDN as a parameter. This causes a command switch from the JCRE context into the currently selected applet's context. When the process is method calls, the VM switches back to the JCRE context. The JCRE sends a response AIDN along with the first command AIDN.

Java™ Card™ Runtime Environment (JCRE) 2.1 Specification

Java™ Card™ Runtime Environment (JCRE) 2.1 Specification

Most method invocations in Java Card technology do not cause a context switch. Context switches only occur during invocation of and return from certain methods, as well as during exception calls from those methods. (See 6.2.1.)

During a context-switching operation, one additional piece of data, indicating the currently active context, is pushed onto the return stack. This context is retained when the method is exited.

Further details of contexts and context switching are provided in later sections of this chapter.

6. Applet Isolation and Object Sharing

Any implementation of the JCRE shall support isolation of contexts and applets. Isolation means that one applet can not access the fields or objects of an applet in another context unless the other applet explicitly provides an interface for access. The JCRE mechanisms for applet isolation and object sharing are detailed in the auxiliary below.

6.1 Applet Firewall

The applet firewall within Java Card technology is runtime-enforced protection and is separate from the Java language technology protections. The Java language protections still apply to Java Card applets. The Java language ensures that strong typing and protection attributes are enforced.

Applet objects are always enclaved in the Java Card VM. They allow the VM to automatically perform additional security checks at runtime.

6.1.1 Contexts and Context Switching

Firewalls essentially partition the Java Card platform's object space into separate protected object spaces called contexts. The firewall is the boundary between one context and another. The JCRE shall allocate and manage an applet's context for each applet that is installed on the card. (This see paragraph 6.1.1.2 below for a discussion of *empty contexts*.)

In addition, the JCRE maintains its own *JCRE context*. This context is much like an applet context, but it has special system privileges so that it can perform operations that are denied to applet contexts. At any point in time, there is only one active context within the VM. (This is called the *currently active context*.) All invocations of methods objects are enclaved in this context against the currently active context. In order to determine if this access is allowed, a Java™ API security interceptor runs at the user's whim and a check is performed.

When certain well-defined conditions are met (including the execution of invoke-type bytecodes as described in paragraph 6.2.8), the VM performs a context switch. The previous context is pushed onto an internal VM stack, a new context becomes the currently active context, and the invoked method executes in this new context. (Upon exit from this installed the VM performs a returning context switch. The original context (of the caller of the method) is popped from the stack and is restored as the currently active context. Context switches can be nested. The maximum depth depends on the amount of VM stack space available.)

6.1.1.1 Group Contexts

Usually, each instance of a Java Card applet defines a separate context. But with Java Card 2.1 technology, the concept of *group contexts* is introduced. It allows that one applet is contained in a single Java package, they share the same context. Additionally, all instances of the same applet class share the same context. In other words, there is no firewall between two applet instances in a group context.

The discussion of contexts and context switching above in section 6.1.1 assumes that each applet instance is associated with a separate context. In Java Card 2.1 technology, contexts are compared to either the file ID, and its instance ID is pushed onto the stack. Additionally, this happens not only when the context switches, but also when context switching occurs on an object owned by one applet instance to an object owned by another instance within the same package.

6.1.2 Object Ownership

When a new object is created, it is associated with the currently active context. (That the object is owned by the applet instance within its currently active context *unless the object is instantiated*. An object is owned by its applet instance, or by the JCRE.)

6.1.3 Object Access

In general, an object can only be accessed by its creating context, that is, when the owning context is the currently active context. The firewall prevents an object from being accessed by another applet in a different context.

In implementation terms, each time an object is accessed, the object's owner context is compared to the currently active context. If these do not match, the access is not performed and a security exception is thrown.

An object is accessed when one of the following bytecodes is executed using the subject's reference:

- `getfield`, `putfield`, `invokevirtual`, `invokeinterface`,
- `attnow`, `opload`, `opstore`, `arraylength`, `checkcast`, `instanceof`,
- `dup` refers to the various types of copy bytecodes, such as `badup`, `bastore`, etc.

This file handles any special or optimized forms of these bytecodes implemented in the Java Card VM, such as `getfield`, `putfield`, `attnow`, etc.

6.1.4 Firewall Protection

The Java Card firewall provides protection against the most frequently anticipated security concerns: de-serializer attacks and denial of service attacks that might allow sensitive data to be "taken over" in another applet. An applet may be able to obtain an object reference from a publicly accessible location, but if the object is owned by a different applet, the firewall ensures security.

EV 263 600 854 US

Docket No. UNP3709CN

Java Card™ Runtime Environment (JCER) 2.1 Specification

The firewall also provides protection against incorrect code. If incorrect code is loaded onto a card, the firewall still protects objects from being accessed by this code.

The Java Card API 2.1 specifies the basic minimum protection requirements of contexts and firewalls because these features should be supported in ways that are transparent to the applet developer. Developers shall be aware of the location of objects, API, and exception related to the firewall.

ICER implementations are free to implement additional security mechanisms beyond those of the standard firewall, as long as these mechanisms are transparent to applets and do not change the externally visible operation of the VBI.

6.1.5 Static Fields and Methods

It should also be noted that classes are not owned by contexts. There is one running context object that can be performed when a class static field is accessed. Native is given a context switch when a static method is invoked. (Similarly, a context switch causes a context switch.)

Public static fields and public static methods are accessible from any context: static methods private to the same context as their caller.

Objects referenced in static fields are from regular objects. They are owned by whatever context created them and standard firewall access rules apply. If it is necessary to share them across multiple applet contexts, then these objects need to be *Show public interface objects* (SPO). (See paragraph 6.2.4 below.)

Of course, the conventional Java technology guarantees are still enforced for static fields and methods. In addition, when objects are instantiated, the instance's static field values attempt to take on an initial static field or, if such is permitted, instantiation and application of static linkage are beyond the scope of this specification.

6.1.5.1 Optional static access checks

The JCER may perform optional runtime checks that are redundant with the constraints enforced by a verifier. A Java Card VM may detect when code violates fundamental language restrictions, such as involving a private method in another class, and report or terminate code that violates.

6.2 Object Access Across Contexts

To enable objects to interact with each other and with the JCER, some well-defined yet laxer mechanisms are provided so one context can access an object belonging to another context.

These mechanisms are provided in the Java Card API 2.1 and are discussed in the following sections:

- **ICER Entry Point Objects**
 - Global Array
 - ICER Privilege
 - Shareable Interfaces

Java Card™ Runtime Environment (JCER) 2.1 Specification**6.2.1 JCER Entry Point Objects**

Secure computer systems shall have a way for them privileged user processes that are restricted to a subset of resources to request system services performed by privileged "system" entities.

In the Java Card API 2.1, this is accomplished using **JCER Entry Point Objects**. These are objects owned by the JCER context, but they have been flagged as containing entry point methods.

The firewall protects these objects from access by applets. The entry point designation allows the methods of these objects to be invoked from any context. When that occurs, a context switch to the JCER context is performed. These methods are the gateways through which applets request privileged JCER system services.

There are two categories of JCER Entry Point Objects :

Temporary JCER Entry Point Objects

Like all JCER Entry Point Objects, methods of temporary JCER Entry Point Objects can be invoked from any applet context. However, references to these objects cannot be stored in class variables, instance variables, or array components. The JCER does not cache attempts to store references to these objects as part of the firewall functionality to prevent unauthorized reuse.

The AID object and all ICER owned exceptions objects are examples of temporary JCER Entry Point Objects.

Permanent JCER Entry Point Objects

Like all JCER Entry Point Objects, methods of permanent JCER Entry Point Objects can be invoked from any applet context. Additionally, references to these objects can be stored and safely re-used.

JCER owned AID instances are examples of permanent JCER Entry Point Objects.

The JCER is responsible for:

- Determining which privileged services are provided to applets.
- Defining contexts containing the entry point methods for these services.
- Creating one or more object instances of these contexts.
- Designating them instances as JCER Entry Point Objects.
- Designating JCER Entry Point Objects as temporary or permanent.
- Making references to these objects available to applets as needed.

Note — Only the methods of these objects are accessible through the firewall. The fields of these objects are protected by the firewall and can only be accessed by the JCER context.

Only the JCER itself can designate Entry Point Objects and wherein they are temporary or permanent. JCER implementations are responsible for implementing the mechanism by which JCER Entry Point Objects are designated and how they become temporary or permanent.

6.2.2 Global Arrays

The global nature of some objects requires that they be accessible from any applet context. This firewall provides ordinarily persistent objects from being used in a flexible manner. The Java Card VM allows an object to be designated as global.

All global arrays and temporary global entry objects. These objects are owned by the JCER context, but can be accessed from any applet context. However, references to these objects cannot be stored in class variables.

EV 263 600 854 US

Docket No.SUN-P3709CNT

Java™ Card™ Runtime Environment (JCRE) 2.1 Specification

instance variables of array components. The API's details and features correspond to share references to these objects as part of the firewall functionality to prevent unauthorized reuse.

For added security, only arrays can be designated as global and only the JCRE itself can designate global arrays. Because applets cannot create arrays, no API methods are provided for manipulating the mechanisms by which global arrays are designated.

At this time of publication of this specification, the only global arrays supported in the Java Card API 2.1 are the API buffers and the byte array input port another (port 0x0) to the applet's local memory.

Note — *Warning*! All of this global arrays, the API specifies that the API buffer is cleared to zero octets whenever an applet is activated, before the Card accepts a new API command. This is to prevent an applet's potentially sensitive data from being "leaked" to another applet via the global API buffers. The API buffer can be accessed from a shared interface object context and is available for passing data between applet contexts. The applet is responsible for protecting secret data that may be accessed from the API buffer.

6.2.3 JCRE Privileges

Because it is the "system" context, this JCRE context has a special privilege. It can invoke a method of any object on the card. For example, assume that object X is owned by applet A. Normally, only context A can access the fields and methods of X. But the JCRE context is allowed to invoke any of the methods of X. During such an invocation, a context switch occurs from the JCRE context to the applet context that owns X.

Note — The JCRE can access both methods and fields of X. Method access is the mechanism by which the JCRE exists as an applet context. Although this JCRE could invoke any method through the firewall, it shall only invoke the newest, previously declared, declared, and get shareable fields and methods (see 6.2.7.1) methods defined in the applet class.

The JCRE context is the currently active context when the VM begins running after a card reset. The JCRE context is the "Card" context and is always claim the currently active context or the previous context saved on the stack.

6.2.4 Shareable Interfaces

Shareable interfaces are a new feature in the Java Card API 2.1 to enable applet interaction. A shareable interface defines a set of shareable interface methods. These interface methods can be invoked from one applet context even if the object implementing them is owned by another applet context.

In this specification, an object instance of a class implementing a shareable interface is called a *shareable interface object (SIO)*.

To this existing context, the SIO is a normal object whose fields and methods can be accessed. To any other context, the SIO is an instance of the shareable interface methods, and only the methods defined in the shareable interface are accessible. All shareable fields and methods of the SIO are protected by the firewall.

Shareable interfaces provide a secure mechanism for inter-applet communication, as follows:

- To make an object available to another applet, applet A first defines a shareable interface, **SI**. A shareable interface extends the interface **IShareable**, **IShareable**, **IShareable**, **IShareable**. The methods defined in the shareable interface, **SI**, represent the services that applet A makes available to other applets.
- Applet A then defines a class **C** that implements the shareable interface **SI**. **C** implements the methods defined in **SI**. **C** may also define other methods and fields, but these are protected by the applet firewall. Only those methods defined in **SI** are accessible to other applets.

Copyright © December 14, 1998 Sun Microsystems, Inc. 6-6

Java™ Card™ Runtime Environment (JCRE) 2.1 Specification

- Applet A creates an object instance **O** of class **C**. **O** belongs to applet A, and the firewall allows A to access any of the fields and methods of **O**.
- To access applet A's object **O**, applet B creates an object reference **SI0** of type **SI**.
- Applet B invokes a special method (`getShareableObject(SI0)`) to request a shared interface object reference from applet A.
- Applet A receives the request and the **IID** of the request (1) via `Applet.getShareableObject(IID)`, and determines whether or not it will share object (1) with applet B.
- If applet A agrees to share with applet B, **A** responds to the request with a reference to **O**. This reference is cast to type **IShareable**, so that none of the fields or methods of **O** are visible.
- Applet B receives the object reference from applet A, casts it to type **SI**, and stores it in object reference **SI0**. Even though **SI0** actually refers to A's object **O**, **SI0** is of type **SI**. Only the shareable interface methods defined in **SI** are visible to **B**. The firewall prevents the other fields and methods of **O** from being accessed by **B**.
- Applet B can request services from applet A by invoking one of the shareable interface methods of **SI0**. During the invocation the Java Card VM performs a context switch. The original (active) active context (2) is saved on a stack and the context of the request (3) (of the actual object (4) because the new (call) active context, **A**'s implementation of the shareable interface method (SI method) executes in **A**'s context.
- The SI method can find out the **IID** of the client (B) via the `getShareableObject(IID)` method. This is described in paragraph 6.2.5. The method determines whether or not it will perform the service for applet B.
- Because of the context switch, the firewall allows the **SI** method to access all the fields and methods of object **O** and any other object owned by **A**. At the same time, the firewall prevents the methods of **SI** from accessing those shared objects owned by **B**.
- The **SI** method can access the parameters passed by **B** and can provide a return value to **B**.
- During the return, the Java Card VM performs a returning context switch. The original (active) active context (2) is restored from the stack, and again becomes the current context.
- Because of the context switch, the firewall allows the **SI** method to access any of its objects and prevents **SI** from accessing non-shared objects owned by **A**.

6.2.5 Determining the Previous Context

When an applet calls `JCRE.setPreviousContext(IID)`, the JCRE shall return the instance **IID** of the previous context active at the time of the last context switch.

6.2.5.1 The JCRE Context

The JCRE context does not have an **IID**. If an applet calls the `getPreviousContext(IID)` method when the applet context was enclaved directly from the JCRE context, this method returns **IID**.

If the applet calls `getPreviousContext(IID)` from a method that may be accessed either from within the applet itself or when accessed via a shareable interface from an external applet, it shall check for a **context** before performing calling **IID** authentication.

Java Card™ Runtime Environment (JCRE) 2.1 Specification

6.2.6 Shareable Interface Details

A shareable interface is simply one that extends (either directly or indirectly) the `java.lang.Interface` interface. Shareable. This shareable interface is similar in concept to the `shareable` interface used by the RMI facility, in which calls to the interface methods take place across a distributed boundary.

6.2.6.1 The Java Card Shareable Interface

Interfaces extending the `Shareable` interface have this special property: calls to this interface methods take place across Java Card's applet firewall boundary via a `shareable` socket.

The Shareable interface is able to identify all shareable objects. Any object that needs to be shared through the applet firewall shall directly or indirectly implement this interface. Only those objects specified in a shareable interface are available through the firewall.

Implementation classes can implement any number of shareable interfaces and can extend other shareable implementation classes.

Like any Java platform interface, a shareable interface defines a set of service methods. A service provider class declares that it "implements" the shareable interface and provides implementations for each of the service methods of the interface. A service client class accesses the service by obtaining an object reference, casting it to the shareable interface type if necessary, and invoking the service methods of the interface.

The shareable interface within the Java Card technology shall have the following properties:

- When a method in a shareable interface is invoked, a mutual switch occurs in the context of the object's context.
- When the method exits, the context of the table is restored.
- Executing `finalize` is guaranteed to allow the currently active object to correctly restore during the `finalize` function finishing (that occurs as an exception is thrown).

6.2.7 Obtaining Shareable Interface Objects

Inter-applet communication is accomplished when a client applet invokes a shareable interface method of a SIO belonging to a server applet. In order for this to work, there must be a way for the client applet to obtain the SIO from the server applet in the first place. The JCRE provides a mechanism to make this possible. The `Applet` class and the `JCRE` system class provide methods to enable a client to request services from the server.

6.2.7.1 The Method `Applet.getShareableInterfaceObject`

This method is implemented by the server applet instance. It shall be called by the JCRE to negotiate between a client applet that requests to use an object belonging to another applet, and the server applet that makes its objects available for sharing.

The default behavior shall return null, which indicates that an applet does not participate in inter-applet communication.

A server applet shall be invoked from another applet block to override this method. This method should respond to the client applet and the parameter. If the client makes a call to one of the expected ALIDs, the method should return null. Similarly, if the parameter is not recognized or it is not allowed for the method, it should return null.

Java Card™ Runtime Environment (JCER) 2.1 Specification

of certain IP, then the method also should return null. Otherwise, the applet should return an SIO of the shareable interface type that the client has requested.

The server applet need respond with this same SIO to all clients. The server can support multiple types of shared interface for different purposes and use client ALID and parameter to determine which kind of SIO to return to the client.

6.2.7.2 The Method `JCRE.getAppletShareableInterfaceObject`

The JCRE class contains the method `getAppletShareableInterfaceObject`, which is invoked by a client applet to communicate with a server applet.

The JCRE shall implement this method to behave as follows:

1. The JCRE searches its internal applet table for one with serverALD. If not found, null is returned.
2. The JCRE looks for this applet's `getShareableInterfaceObject` method, passing the clientALD, the client, and the parameter.
3. A context which occurs to the server applet, and its implementation of `getShareableInterfaceObject` proceeds as described in the previous section. The server applet returns a SIO (or null).
4. `getAppletShareableInterfaceObject` method returns the same SIO (or null) to its caller.

For enhanced security, the implementation shall make it impossible for the client to tell which of the following conditions caused a null value to be returned:

- The serverALD was not found.
- The server applet does not recognize the clientALD and returns `getShareableInterfaceObject`.
- The server applet was not communicating with this client.
- The server applet was not communicating with this client as specified by the parameter.

6.2.8 Classes and Object Access Behavior

A static class field is guaranteed when one of the following Java bytecode is executed:

`getStatic`, `putStatic`

An object is guaranteed when one of the following Java bytecode is executed using the object's reference:

`putField`, `invokeStatic`, `invokeInterface`, `absorb`, `create`, `createLocal`, `arrayLength`, `cancelLocal`, `cancel`.

<?> refers to the various types of array objects, such as `byte[]`, `short[]`, `int[]`, etc.

This class includes any special or optimized forms of these bytecode that may be implemented in the Card VM, such as `get Field`, `get Field A`, `get Field B`, `get Field C`, `absorb`, `create`, `cancel`.

Prior to performing the work of this bytecode as specified by the Java Card VM, the Java Card VM will perform a check on the referenced object. If object is denied, then a `SecurityException` is thrown.

The access checks performed by the Java Card VM depend on the type and context of the referenced object, the bytecode, and the currently active method. They are described in the following sections.

Java™ Card™ Runtime Environment (JCER) 2.1 Specification

6.2.8.1 Accessing Static Class Fields

By code:

getStatic, putStatic

If the JCER is the currently active context, then access is allowed.

- Otherwise, if the bytecode is parameter and the field being stored is a reference type and the reference being stored is a reference to a temporary JCER Entry Point Object or a global array then access is denied.
- Otherwise, access is allowed.

6.2.8.2 Accessing Array Objects

By code:

getArray, setArray, getLength, getElement, insertElement

If the JCER is the currently active context, then access is allowed.

- Otherwise, if the bytecode is a static field and the component being stored is a reference type and the instance being stored is a reference to a temporary JCER Entry Point Object or a global array then access is denied.
- Otherwise, if the array is owned by the currently active context, then access is allowed.
- Otherwise, if the array is designated global, then access is allowed.
- Otherwise, access is denied.

6.2.8.3 Accessing Class Instance Object Fields

By code:

getField, getField

If the JCER is the currently active context, then access is allowed.

- Otherwise, if the bytecode is a static field and the field being stored is a reference type and the reference being stored is a reference to a temporary JCER Entry Point Object or a global array then access is denied.
- Otherwise, if the object is owned by the currently active context, then access is allowed.
- Otherwise, access is denied.

6.2.8.4 Accessing Class Instance Object Methods

By code:

invokeVirtual

- If the object is owned by the currently active context, then access is allowed. Context is switched to the object owner's context.
- Otherwise, if the object is delegated a JCER Entry Point Object, then access is allowed. Context is switched to the object owner's context (still in JCER).

Java™ Card™ Runtime Environment (JCER) 2.1 Specification

■ Otherwise, if ACBS is the currently active context, then access is allowed. Context is switched to the object owner's context.

■ Otherwise, access is denied.

6.2.8.5 Accessing Standard Interface Methods

By code:

invokeInterface

- If the object is owned by the currently active context, then access is allowed.
- Otherwise, if the ACBS is the currently active context, then access is allowed. Context is switched to the object owner's context.
- Otherwise, access is denied.

6.2.8.6 Accessing Shareable Interface Methods

By code:

invokeInterface

- If the object is owned by the currently active context, then access is allowed.
- Otherwise, if the object's class implements a shareable interface, and if the interface being invoked extends the Shareable interface, then access is allowed. Context is switched to the object owner's context.
- Otherwise, if the ACBS is the currently active context, then access is allowed. Context is switched to the object owner's context.
- Otherwise, access is denied.

6.2.8.7 Throwing Exception Objects

By code:

at throw

- If the object is owned by the currently active context, then access is allowed.
- Otherwise, if the object is designated a JCER Early Point Object, then access is allowed.
- Otherwise, if no JCER is the currently active context, then access is allowed.
- Otherwise, access is denied.

6.2.8.8 Accessing Class Instance Objects

By code:

checkcast, instanceof

- If the object is owned by the currently active context, then access is allowed.
- Otherwise, if JCER is the currently active context, then access is allowed.
- Otherwise, if the object is designated a JCER Early Point Object, then access is allowed.
- Otherwise, if the JCER is the currently active context, then access is allowed.
- Otherwise, access is denied.

6.2.8.9 Accessing Standard Interfaces

By code:

checkcast, instanceof

- If the object is owned by the currently active context, then access is allowed.
- Otherwise, if the JCER is the currently active context, then access is allowed.
- Otherwise, access is denied.

6.2.8.10 Accessing Shareable Interfaces

By code:

checkcast, instanceof

- If the object is owned by the currently active context, then access is allowed.
- Otherwise, if the object's class implements a shareable interface, and if the object is being cast into (classcast) or is an instance of (instanceof) an interface that extends the shareable interface, then access is allowed.
- Otherwise, if the JCER is the currently active context, then access is allowed.
- Otherwise, access is denied.

object is **conditionally updated**. The field or array component **appears** to be updated—traversing the **ObjectArray** compound **list** yields the **modified value**—but the update is not yet **committed**.

When the **apply** call **executes**, **create** a **Transaction**. All conditional updates are **committed** to persistent storage. If power is lost or if some other system failure occurs **fails** to the **completion** of **create**, the **ObjectArray** **consistently** reflects all conditionally updated fields or array components are restored to their previous values. If the update **occurred** as a **logical problem** or **decided** to **cancel** the transaction, it can programmatically undo conditional updates by calling **cancel**, **abort** or **Transaction::cancel**.

7. Transactions and Atomicity

A **transaction** is a logical set of updates of persistent data. For example, instantiating some amount of memory for one account or **balance** is a **single transaction**. It is **isolation** for transaction to be atomic; either all of the data fields are updated, or none are. The JCRE provides robust support for **atomic transactions**, so that read data is restored to its original pre-transaction state if the transaction does not complete normally. This **transaction** protects against events such as power loss in the middle of a transaction, and **multiple program errors** that might cause data corruption should all steps of a transaction not complete normally.

7.3 Transaction Duration

A **transaction** **begins** and **ends** when the JCRE **receives** group **automatic** update **return** from the **Object's** **setObject**, **create**, **process**, or **cancel** methods. This is true whether a transaction ends naturally, with an object's **call** to **cancel**, **process**, or, with an **abortion** of the transaction, **cancel** **programmatically** by the **Object**, as by default by the JCRE. For more details on **transaction abortion**, refer to [page 7-6](#).

Transaction duration is the life of a transaction between the **call** to **Object::start** and **Object::cancel**, and **either** a **call** to **cancel** **Transaction** or an **abortion** of the transaction.

7.1 Atomicity

Atomicity **declares** that the card handles the **creation** of persistent storage after a step. Failure of fatal exception during an update of a single object or class field or array component. If power is lost during the update, this **apply** **dispatcher** shall be able to rely on what the field or array component **contains** when power is restored.

The Java™ Card platform guarantees that any update to a single persistent object or class field will be atomic. In addition, the Java™ Card platform guarantees **atomic** **copy** **operations** **local** **atomically** for persistent arrays. That is, if two smart card code powers during the update of a data element (field) in an object/class or **copy** **operation** of an array that shall be present and correct (CAS) **atomic**, that data element will be restored to its previous values.

Smart card code guarantees atomality for block updates of multiple data elements. For example, **two** **atomicity** of the **set** or **copy** method guarantees that either all bytes are correctly copied or else the destination array is restored to its previous byte values.

An **apply** might not require atomality for **array update**. The **set** or **arrayCopy** **methods** **is** **provided** for this purpose. It does not use the **transaction** **context** **buffer** even when called with a **transaction** in progress.

7.2 Transactions

An **apply** **method** **attempts** to atomically update several **different** **fields** or **array components** in several **different** **objects**. Either all updates take place **atomically** and **consistently**, or else all **fields/components** are restored to their previous values.

The Java™ Card platform **supports** a **transactional model** in which an **apply** **can** **designate** the **beginning** of an **atomic** **set** of **updates** with a **call** to the **Object::start** **method**. Each **object** **updates** after this **beginning** **are** **atomic** **and** **consistent** with the **Object::cancel** **method**.

If power is lost (reset) or the card is reset or some other system failure occurs while a transaction is in progress, then the JCRE shall restore to which previous values all fields and array components **conditionally update** since the previous call to **Object::start**, **Object::cancel**.

This action is performed automatically by the JCRE when it **cancelizes** the card after encounters from the power loss, reset, or failure. The JCRE determines which of those objects (if any) were conditionally updated and restores them.

Note – Object update used by instances created during the transaction that failed due to power loss or card reset can be recovered by the JCRE.

Java™ Card™ Business Requirements (JCER) 2.1 Specification

7.6 Aborting a Transaction

Transactions can be aborted either by an applet or by the JCER.

7.6.1 Programmatic Abortion

If an applet encounters an internal problem or decides to cancel the transaction, it can programmatically make conditional updates by calling JCSystem.abortTransaction. If this method is called, all conditionally updated fields and every component, since the previous call to `setTransient`, in any transaction are restored to their previous values, and the JCER sets a `transient` value to `0`.

7.6.2 Abortion by the JCER

If an applet returns from `select`, `doSelect`, `process`, or `Install` methods with a transaction by program, the JCER automatically aborts the transaction. If a return from any of `select`, `doSelect`, `process` or `Install` methods occurs with a transaction in progress, the JCER acts as if an exception has thrown.

7.6.3 Cleanup Responsibilities of the JCER

Object instances created during the transaction that is being aborted can be deleted only if all references to these objects can be located and converted into null. The JCER itself ensures that references to objects created during the aborted transaction are equivalent to a null reference.

7.7 Transient Objects

Only updates to persistent object participants in the transaction. Updates to transient objects are never transient, regardless of whether or not they were "pinned" to a transaction.

7.8 Commit Capacity

Since platform resources are limited, the number of bytes of conditionally updated data that can be accumulated during a transaction is limited. The Java Card technology provides methods for determining how much transient capacity is available on the target application. The transient capacity represents an upper bound on the number of conditionally updated bytes available. The actual number of conditionally updated bytes available may be lower due to management overhead.

An exception is thrown if the transient capacity is exceeded during a transaction.

8. API Topics

8.1.1.1 Captured Transfers with no chaining
 When the API chaining mode of output transfer is requested by the applet by calling the `SetOutgoingTransferInChaining` method, the following protocol sequence shall be followed.

Note -- when the no chaining mode is used, call to `GetOutgoingTransferIn` method shall throw an APDU exception with reason code `ILLEGAL_USE`.

Notation

`Lo = CAD expected length.`

`Le = Applet response length set via SetOutputLength method.`

`<IN3> = the protocol byte equal to the incoming transfer IN3 byte, which indicates that all data bytes will be transferred as is.`

`<=IN3> = the protocol byte that is the complement of the incoming header IN3 byte, which indicates that 1 data byte being transferred as is.`

`<SW1,SW2> = the response status bytes as in ISO7816-4.`

ISO 7816-4 CASE 2

`Lo = Lo`

1. The card sends `Le` bytes of output data using the standard `T=0 <IN3> or <=IN3>` procedure byte mechanism.
2. The card sends `<SW1,SW2>` completion status on completion of the Applet. procedure

`Le = Lo`

1. The card sends `Le` bytes of output data using the standard `T=0 <IN3> or <=IN3>` procedure byte mechanism.
2. The CAD sends `Le` bytes of response using the standard `T=0 <IN3> or <=IN3>` procedure byte mechanism.
3. The card sends `<SW1,SW2>` completion status on completion of the Applet. procedure
4. The card sends `<SW1,SW2>` completion status on completion of the Applet. procedure

Resource Use within the API
 Unless specifically called out in the Java Card 2.1 API Specification, the implementation shall support the following of API features initialized, even when the owner of the object instance is not the currently selected applet. In other words, unless specifically called out, the implementation shall not use resources such as transient objects of `CLEAR_ON_DESELECT` type.

Exceptions thrown by API classes

All exception objects thrown by the API implementation shall be temporary JCME Event Point Objects. Temporary JCME Event Point Objects cannot be stored in class variables, instance variables or array components. (See 8.2.1)

8.1 The APDU Classes

The APDU class encapsulates access to the ISO 7816-4 based APDU services via the serial line. The APDU Class is designed to be independent of the underlying APDU transport protocols or busses.

The JCER may support `T=0` or `T=1` transport protocols or both.

8.1.1 T=0 specifics for outgoing data transfers

For compatibility with legacy JCDEnterprise, this section supports APDU chaining to implement the APDU Class. All APDU selection via the `GetOutgoingTransferIn` method.

1. The card sends `Le` bytes of output data using the standard `T=0 <IN3> or <=IN3>` procedure byte mechanism.
2. The card sends `<SW1,SW2>` completion status bytes.
3. The CAD sends `Le` bytes of response using the standard `T=0 <IN3> or <=IN3>` procedure byte mechanism.
4. The card sends `<SW1,SW2>` completion status on completion of the Applet. procedure

Java™ Card™ Radius Environment (JCRE) 2.1 Specification

3. Repeat steps 2-4 as necessary to send the remaining output data bytes (L4) as required.
6. The card reads <SW1,SW2> completion status to continuation of the APDU st. process as needed.

ISO 7816-4 CASE 4

In Case 4, L4 is determined after the following initial readings:

1. The card sends <0x00, L4> bytes by code.
2. The CAD sends GET RESPONSE command with Lc <= L4.

The rest of the protocol sequence is identical to CASE 3 described above.

If the output bytes partly and sends less than L4 bytes, set Lc may be issued instead to tell out the length of the bytes expected by the CAD.

8.1.1.2 Regular Output Transfers

When the no chaining mode of output transfers is not requested by the applet (thus Lc, the `setOutputChaining` method is used), the following protocol sequence shall be followed:

Any ISO-7816-3/M message L=0 protocol transfer sequence may be used.

Note — The `write` function method may be invoked by the applet between successive calls to `sendByCard` or a card `read` or `long` method. The `write` function can be used that ends up S-block command with `AT&T` response of IMF wait, which is equivalent to a request of 1 additional work waiting time in T=0 mode. (See ISO 7816-3).

8.1.2 T=1 specifics for outgoing data transfers

8.1.2.1 Centralized transfers with no chaining

When the no chaining mode of output transfers is requested by the applet by calling the `setOutputChaining` method, the following protocol sequence shall be followed:

Notation

$L4 = \text{APDU response length, set via } \text{setOutputChaining} \text{ method.}$

Java™ Card™ Radius Environment (JCRE) 2.1 Specification

This transport protocol sequence shall not use block chaining. Specifically, the N-bit (more data bits) shall be sent in the PCH of the L-block during the transfer (ISO 7816-3). In other words, the entire outputting data (L-block type) shall be transferred in one L-block.

(If the applet sends early and sends less than L4 bytes, zeros shall be sent instead of L4 and the remaining length of the block.)

Note — When the no chaining mode is used, calls to `write` that generate a method shall do on an APDU reception with receive code <1,1,0,0,1,0>, USE.

8.1.2.2 Regular Output Transfers

When the no chaining mode of output transfers is not requested by the applet (i.e. the `setOutputChaining` method is used), the following protocol sequence shall be followed:

Any ISO-7816-3/M sequence T=1 protocol transfer sequence may be used.

Note — The `write` function method may be invoked by the applet between successive calls to `sendByCard` or a card `read` or `long` method. The `write` function can be used that ends up S-block command with `AT&T` response of IMF wait, which is equivalent to a request of 1 additional work waiting time in T=0 mode. (See ISO 7816-3).

8.2 The security and crypto packages

The `get` instance method in the following classes return an implementation instance in the context of the calling applet of the requested algorithm:

- javocard.security.KeyPairGenerator
- javocard.security.Signature
- javocard.security.RandomData
- javocard.security.Crypto
- javocard.Cipher

An implementation of the JCRE may implement 0 or more of the algorithms listed in the API. When an algorithm that is not implemented is requested this method shall throw a `CryptoException` with cause code `0x8000000000000001`.

Implementation of the above classes shall extend the corresponding base class and implement all `get` methods. All data allocation associated with the implementation instance shall be performed in the `get` instance construction to ensure that any lack of required resources can be diagnosed early during the `get` instance of the applet.

Similarly, the `buildKey` method of the `JavaCard`, `Security` and `Signature` classes return an implementation instance of the requested key type. The JCRE may implement 0 or more types of keys. When a key type that is not implemented is requested this method shall throw a `CryptoException` with cause code `0x8000000000000002`.

Java Card and Runtime Environment (JCRE) 2.1 Specification

Implementation of key types shall implement the associated interface. All data alterations associated with the key implementation interface shall be performed at the time of instance construction to ensure that any lack of required resources can be flagged early during the initialization of the object.

8.3 JCSystem Class

In Java Card 2.1, the getVersion method shall return (short) 0x0201.

9. Virtual Machine Topics

The topics in this chapter detail virtual machine specifics.

9.1 Resource Failures

A lack of free memory (such as heap space) which is recoverable should result in a `SystemException` with reason code `NO_MEMORY`. The `free()` method in `Activator` can be used to create transient errors through a system-dependent loss of memory.

All other (non-recoverable) virtual machine errors such as stack overflow and result in a `VirtualMachineError`. These conditions shall cause the virtual machine to fail. When such a non-recoverable virtual machine error occurs, an implementation can optionally cause the card to be muted or blocked from further use.

Java™ Card™ Runtime Environment (JCRE) 2.1 Specification

Java™ Card™ Runtime Environment (JCRE) 2.1 Specifications

Obviously, a JCIRB implementer could choose to implement the installers as an API. If so, then the installers might be coded to interact with application files and respond to invocations of the select, process, and deregister methods.

But a JCIRB implementer could also implement the installer in other ways, as long as it provides the SELCTable behavior to the outside world. In this case, the JCIRB implementer has the freedom to provide some other mechanism by which APDUs are delivered to the installer code module.

10. Applet Installer

Applet installation is similar to using Java Card technology as a complete topic. The Java Card API 2.1 is intended to give JCRE implementers as much freedom as possible in fact implementation. However, some basic common specifications are required in order to allow Java Card applets to be installed without knowing the implementation details of a particular installer.

This specification defines the concept of an installer and specifies minimal installation requirements in order to achieve interoperability across a wide range of possible installer implementations.

The Applet Installer is an optional part of the JCRE 2.1 Specification. That is, an implementation of the JCIRB does not necessarily need to include a Java Card Installer. However, if implemented, the installer is required to support the behavior specified in section 9.3.

The mechanisms necessary to install an applet on a smart card using Java Card technology are embodied in an API card component called the *Installer*.

To the CAD the installer appears to be an applet. It has an API, and it becomes the currently selected applet when this API is successfully processed by a SELCT command. Once selected, the installer behaves in much the same way as any other applet:

- It receives all APDUs, but like any other selected applet.
- Its design specification prescribes the various kinds and formats of APDUs that it expects to receive along with the genericities of those commands under various preconditions.
- It processes and responds to all APDUs that it receives. Incorrect APDUs are responded to with an error condition of Java Card.
- When another applet is selected (or when the card is reset or when power is removed from the card), the installer becomes deselected and remains deselected until the next time that it is SELCTed.

10.1.1 Installer Implementation

The installer need not be implemented as an applet on the card. The requirement is only that the installer functionally be SELCTable. The corollary to this requirement is that installer component shall not be able to be invoked when other functionality higher in the class hierarchy of Java Card applets is selected.

Obviously, a JCIRB implementer could choose to implement the installers as an API. If so, then the installers might be coded to interact with application files and respond to invocations of the select, process, and deregister methods.

But a JCIRB implementer could also implement the installer in other ways, as long as it provides the SELCTable behavior to the outside world. In this case, the JCIRB implementer has the freedom to provide some other mechanism by which APDUs are delivered to the installer code module.

10.1.2 Installer API

Because the installer is SELCTable, it shall have no API. JCIRB implementations are free to choose the API by which their installer is selected. Multiple installers may be implemented.

10.1.3 Installer APDUs

The Java Card API 2.1 does not specify any APDUs for the installer. JCIRB implementations are entirely free to choose their own API and communicate to derive their installer to be a result. The model is that the installer on the card is driven by an installation proxy running on the CAD. An installer installation to succeed, this CAD installation proxy must be able to:

- Recognize the card.
- SELCT the installer on the card.
- Define the installation process by sending the appropriate APDUs to the card installer. These APDUs will consist of:
 - Authentication information, to ensure that the installation is authorized.
 - The specific code to be loaded into the card's memory.
 - Linkage information to link the applet code with code already on the card.
 - Instance initialisation parameters to be sent to the applet's *Install* method.

The Java Card API 2.1 does not specify the details of the CAD installation program nor the APDUs passed between it and the installer.

10.1.4 Installer Behavior

JCRE implementers shall also define other behaviors of their installer, including:

- Whether or not installation can be aborted and how this is done.
- What happens if an exception, error, or power & UI occurs during installation.
- What happens if another applet is selected before the installer is finalized with its code.

The JCIRB shall guarantee that an applet will not be deemed successfully installed if:

- the applet's *Install* method throws an exception before successful return from the *applets.request* method. (Refer to paragraph 9.2.)

Java™ Card™ Routine Environment (JCER) 2.1 Specification**10.1.5 Installer Privileges**

Although an installer may be implemented as an applet, an installer will typically require access to features that are not available to "other" applets. For example, depending on the JCER implementation's implementation, the installer will need to:

- Read and write directly to memory bypassing the object's protection and/or standard security.
- Access objects owned by other applets or by the JCER.
- Invoke any security related methods of the JCER.
- Be able to invoke the `Install()` method of a newly installed applet.

Again, it is up to each JCER implementation to determine the features implementation and supply such features. In these JCER implementations as necessary to support their installer, JCER implementations *proactively* make responsible for the security of such features, so that they *proactively* available to normal applets.

10.2 The Newly Installed Applet

There is a single interface between the installer and the applet that is being installed. After the installer has *successfully* prepared the applet for execution (performed steps such as loading and linking), the installer shall invoke the applet's `Install()` method. This method is defined by the `Applet` interface.

The precise mechanism by which an applet's `Install()` method is invoked from the installer is a JCER implementation-defined implementation detail. However, there should be a *standard* switch so that any context-related operations performed by the `Install()` method (such as creating new objects) are done in the context of the new applet and not in the context of the installer. The installer shall also ensure that any objects created during applet class initialization (static) variables are also owned by the context of the new applet.

The installation of an applet is deemed complete if all steps are completed without a error or an exception being thrown, up to and including universal return from executing the `Install()` method or *finished*. At that point, the installed applet will be accessible.

The maximum size of the parameter data is 22 bytes. And for security reasons, the parameter is encoded after the return from the `ARDU` buffer is erased (return from an applet's `process` method.)

10.2.1 Installation Parameters

Other than the maximum size of 22 bytes, the Java Card API 2.1 does not specify anything about the contents of the installation parameter byte array argument. This is fully defined by the applet designer and can be in any format desired. In addition, these installation parameters are intended to be opaque to the installer.

JCER implementations should design their installers so that is a *standard* for an installation program (working in a C API) to specify an arbitrary byte array to be delivered to the installer. This installer simply (research this byte array to the installed applet's `Install()` method in the `ARDU` parameter. A typical implementation might define a JCER implementation-specific API command that has the sequence "call the applet's `Install()` method" a installing this command will be accompanying byte array."

11. API Constants

Some of the API classes don't have values specified for fields contained in the Java `Const` API 2.1 Reference. In constant values are not specified consistently by implementations of this API 2.1 Specification, usually when interoperability is forgotten. This chapter provides the required values for constants that are not specified in the Java Const API 2.1 Reference.

Java™ Card™ Runtime Environment (JCRE) 2.1 Specification

Digitized by Google

Digitized by Google

Glossary

API is an acronym for Application Programming Interface as defined in ISO 7016-4.

APDU is an acronym for Application Protocol Data Unit as defined in ISO 7016-4.

API is an acronym for Application Programming Interface. The API defines coding conventions by which an application program accesses the operating system and other services.

Apple is the trademark of this document for Java Card Appliance, which is the basic unit of application, security, functionality, and security in Java Card technology.

Apple developer refers to a person creating a Java Card Appliance using the Java Card technology specifications.

Apple friend is the application in the Java Card technology by which the VM prevents an application from sending unauthorized messages to objects created by other application contexts or the JCRC context, and reports or other who addressed the violation.

Atomic operation is an operation that either completes in its entirety (if the operation succeeds) or no part of its operation completes at all (if the operation fails).

Attestation refers to whether a particular operation is atomic or not and is necessary for proper data integrity in cases in which power is lost or the card is unexpectedly removed from the CAD.

ATM is an acronym for Answer to Reader. An ATM is a string of bytes sent by the Java Card after a read command.

CAD is an acronym for Card Acceptance Device. The CAD is the device to which the card is inserted.

Card is the physical container from our day to day to store data.

CFCA is the term used to verify the compatibility of the implementation of the Java Card Technology specification. The CFCA uses the JavaCard tool to run the test suite.

Class is the prototype for an object in an object-oriented language. A class may also be considered a set of objects that share a common structure and behavior. The structure of a class is determined by the class variables that represent the state of an object of the class and the behavior to be given by a set of methods associated with the class.

Classes are related in a class hierarchy. One class may be a specialization (a subclass) of another (its super class). It may have references to other classes, and it may use other classes in a client-server relationship. **Client** (See *Applies exclusively*.)

Currently active context The JCRC keeps track of the currently active Java Card Appliance context. When a virtual method is invoked on an object, and a context switch is required and permitted, the currently active

context is changed to correspond to the application context that owns the object. When that method returns, the previous context is restored. Invocations of static methods have no effect on the currently active context. The currently active context and starting status of an object together determine its access to an object in memory.

Currently active object The JCRC keeps track of the currently activated Java Card Appliance. Upon receiving a SELECT command with this object's AID, the JCRC makes this object the currently selected object. The JCRC sends all APDU commands to the currently selected object.

EEPROM is an acronym for Electrically Erasable, Programmable Read Only Memory.

Exception (See *Applies Generally*.)

Framework is the set of classes that implement the API. This includes core and extension packages. Responsibilities include dispensing of APDUs, object selection, managing memory, and handling application.

Garbage collection is the process by which dynamically allocated storage is automatically reclaimed during the execution of a program.

Instance variables, also known as fields, represent a portion of an object's instant state. Each object has its own set of instance variables. Objects of the same class will have the same instance variables, but each object can have different values.

Initialization, in object-oriented programming, means to produce a particular object from its class template. This involves allocation of data structures with the types specified by the template, and initialization of instance variables with either default values or those provided by the class's constructor function.

JAR is an acronym for Java Archive. JAR is a platform-independent file format that combines many files into one.

Java Card Runtime Environment (JCRC) consists of the Java Card Virtual Machine, the Framework, and the associated native methods.

JC1848 is an acronym for the Java Card 2.1 Reference Implementation.

JCRA implementation refers to a person creating a vendor-specific implementation using the Java Card API.

JCVA is an acronym for the Java Card Virtual Machine. The JCVA is the foundation of the JC and portable applications (firmware) and enables secure data sharing.

JDK is an acronym for Java Development Kit. The JDK is a Sun Microsystems, Inc. product that provides the environment required for programming in Java. The JDK is available for a variety of platforms, but notably Sun Solaris and Microsoft Windows.

Method is the name given to a procedure or function, associated with one or more classes in object-oriented languages.

NameSpace is a set of names in which all names are unique.

Object-Oriented is a programming methodology based on the concept of an object, which is a data structure encapsulated with a set of routines, called methods, which operate on the data.

Object, in object-oriented programming, are unique instances of a data structure defined according to the template provided by its class. Each object has its own values for the variables belonging to its class and can respond to the messages (methods) defined by its class.

Copyright © December 14, 1998 Sun Microsystems, Inc.

Java on CAD in Realtime Environment (ICRE) 2.1 Specification

Package is a namespace within the Java programming language and contains classes and interfaces. A package is the smallest unit within the Java programming language.

Persistent object Persistent objects and their values persist from one CAD session to the next. Individually, objects are persistent by default. Persistent objects values are updated atomically using transactions. The term persistent does not mean there is an object-oriented database on the card or that objects are serialized/certified, just that the objects are not lost after the card loses power.

Shareable interface Defines a set of shared interface methods. These interface methods can be leveraged from one application up to the object implementation which is exempt by sending explicit context.

Shareable interface object (SIO) An object that implements this shareable interface.

Transaction is an atomic operation in which the developer defines the range of the operation by indicating to the program both the beginning and end of the transaction.

Transient object. The values of transient objects do not persist between CAD sessions to the card, and are reset to a default value at application startup. Updates to the values of transient objects are not atomic, and are not affected by transactions.

EV 263 600 854 US
Date: 16 December 1998

Docket No.SUN-P3709CNT

Dear Java Card Licensee,

JCRE21-DF2-14DEC98.zip contains a second draft of the Java Card 2.1 Runtime Environment specification, dated December 14, 1998, for Licensee review and comment. We have worked to incorporate and clarify the document based upon the review feedback we've received to date.

Complete contents of the zip archive are as follows:

READ-ME-JCRE21-DF2.txt - This READ ME text file
JCRE21-DF2.pdf - "Java Card Runtime Environment (JCRE) 2.1 Specification" in PDF format
JCRE21-DF2-changebar.pdf - The revised document with change bars from the previous version for ease of review.

Summary of changes:

1. This is now a draft 2 release and will be published on the public web site shortly.
2. New description of temporary JCRE Entry Point Objects has been introduced for purposes of restricting unauthorized access. Firewall chapter 6.2.1.
3. Global arrays now have added security related restrictions similar to temporary JCRE Entry Point objects. Firewall chapter 6.2.2.
4. Detailed descriptions of the bytecodes with respect to storing restrictions for temporary JCRE Entry Point Objects and Global arrays added. Chapter 6.2.8.
5. General statement about JCRE owned exception objects added in chapter 8.
6. Corrected description of Virtual machine resource failures in transient factory methods. Chapter 9.1.

The "Java Card Runtime Environment 2.1 Specification" specifies the minimum behavior and runtime environment for a complete Java Card 2.1 implementation, as referred to by the Java Card API 2.1 and Java Card 2.1 Virtual Machine Specification documents. This specification is required to ensure compatible operation of Java Card applets. The purpose of this specification document is to bring all the JCRE elements together in a concise manner as part of the Java Card 2.1 specification suite.

Please send review comments to <javaoem-javacard@sun.com> or to my address as below. On behalf of the Java Card team, I look forward to hearing from you.

Best,
Godfrey DiGiorgi

Godfrey DiGiorgi - godfrey.digiorgi@eng.sun.com
OEM Licensee Engineering
Sun Microsystems / Java Software
+1 408 343-1506 - FAX +1 408 517-5460